

The Epic Chronicle of Designing Cassini's Titan Flyby Altitudes

David Seal
Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Dr., Pasadena, CA
818-434-7935
seal@jpl.nasa.gov

Molly Bittner
Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Dr., Pasadena, CA
818-354-8815
molly.bittner@jpl.nasa.gov

Abstract— The selection and optimization of Titan flyby altitudes for NASA's Cassini mission at Saturn has traveled a long, fascinating, and often torturous sixteen-year path - starting in 2001, when pre-arrival trajectory design decisions had to be made, through April of 2017 when Cassini's last, and arguably most critical, low flyby takes place. The chronicle of designing and updating the Titan flyby altitudes have twists and turns enough to satisfy a full-length novel or feature film, including: critical design decision-making before arrival with multiple atmospheric models, high uncertainties, and limited data; early flybys that seemed to show trends that weren't there; use and misuse of statistical analysis; unexpected surprises with limited reaction time; navigation of a scientific, engineering, and management community with a wide array of inherent biases; and consideration of a variety of project risk postures in a high-scrutiny, high-impact, high-reward environment.

TABLE OF CONTENTS

1. INTRODUCTION.....	1
2. THE CASSINI ENVIRONMENT PRE-ARRIVAL.....	2
3. TITAN FLYBY PLANNING PRE-ARRIVAL	4
3. EARLY MISSION RESULTS	7
4. THE T16 ANALYSIS	8
5. THE REST OF THE MISSION	9
6. CONCLUSIONS.....	10
ACKNOWLEDGEMENTS.....	11
REFERENCES	11
BIOGRAPHIES	12

1. INTRODUCTION

The story of Titan, to which Cassini has contributed so much since its arrival at Saturn in 2004, begins with its discovery in 1655 by Dutch astronomer Christiaan Huygens, after whom Cassini's Titan atmospheric and surface probe was named. Though Galileo was the first to observe Saturn through a telescope some 45 years earlier, he was unable either to articulate the structure of the rings or to spot Saturn's largest satellite; both were accomplished by Huygens which he documented in his 1655 article *De Saturni Luna Observatio Nova*, or "*A New Observation of Saturn's Moon*", and his renowned *Systema Saturnium* published in 1659. Subsequent Earth-based observations of Titan contributed relatively little save for refining its coarse

orbital geometry, until the 20th century. Hints of an atmosphere around Titan were detected in 1908, when Josep Comas i Solá observed limb darkening. The first conclusive evidence of an atmosphere came from spectroscopic detection of gaseous methane by Gerald Kuiper in 1944. Pioneer 11 was the first spacecraft to visit the Saturnian system in September of 1979, but at an altitude of 363,000 km from Titan, so it was not close enough to make useful observations of the moon or its atmosphere. Voyager 1, in November of 1980, was directed to the first close flyby at 3,900 km from the moon's surface, and collected the bulk of the useful data used for Cassini mission planning before its arrival. Voyager 1 conducted a variety of observations of Titan, including pictures (see Figure 1) and an ultraviolet solar occultation probing the depth and density of its atmosphere.




Figure 1. Titan as seen from Voyager 1; note the north/south hemispherical difference in brightness

In the images of Titan, Figure 1 in particular, a hemispherical difference was noted in the brightness of Titan's atmosphere; HST observations a decade later suggested that the brightness of the hemispheres had

switched, offering tantalizing hints at seasonal change – at least in the atmosphere. Cassini has since discovered that Titan does indeed see seasonal effects, including the appearance and disappearance of islands in the methane lakes on its surface. This is driven not only by Saturn’s (and Titan’s) axial tilt with respect to the ecliptic of 27° but its moderate orbital eccentricity of 0.056; its solar distance varies from 9 AU at perihelion to 10.1 AU at aphelion. Voyager 2 was the last spacecraft to visit Saturn before Cassini and passed even farther away from Titan than Pioneer 11 at an altitude of over 660,000 km in August 1981 – again not close enough to collect additional data.

It is interesting to note, as an aside, that an alternate mission plan for Voyager 1 would have put it on a trajectory to encounter Pluto in 1986 – nearly three decades before New Horizons eventually reached the dwarf planet – but at the cost of the close Titan flyby. It is hard to find fault in the Voyager project’s decision to select Titan based on what was known at the time and the “bird in the hand” perspective, given the long flight time from Saturn to Pluto. (We of course now know that Voyager 1 would have survived the journey.) Titan was such a valuable scientific target, in fact, that had Voyager 1 failed in its observations, Voyager 2 would have been directed to make a repeat attempt, which would have sacrificed its subsequent “grand tour” exploration of Uranus and Neptune in 1986 and 1989.

We now know that Titan is the only satellite in the Solar System with an appreciable atmosphere. It has an atmosphere which is composed mostly of nitrogen, like Earth’s, is 45% more dense at the surface than Earth’s, and masses 20% more than Earth’s in total. This atmosphere contributes to a methane cycle that closely resembles Earth’s water cycle, and makes the satellite one of the most important bodies of scientific study in the Solar System – a clear priority for Cassini, and one worthy of its own probe.

The Cassini mission is a nearly four billion dollar international cooperative effort of NASA, the European Space Agency (ESA), and the Italian Space Agency (ASI), to explore the planet Saturn and its environment. The mission has been a keystone of NASA’s modern program of exploration of the solar system. Cassini launched from Cape Canaveral in October of 1997 and after a seven year cruise, conducted a four-year primary mission scientific exploration of Saturn, its atmosphere, moons, rings and magnetosphere from June 2004 to June 2008. The first mission extension, the Cassini Equinox Mission, completed in September 2010. The second extension, the Cassini Solstice Mission, ends in September 2017 with spacecraft disposal into Saturn’s atmosphere and will complete a period of over 13 years of in-situ study. The Cassini spacecraft consists of the orbiter with twelve instruments and the detachable Huygens Probe, which descended to the surface of Titan successfully in January 2005.

Not only has Titan been a high priority for the Cassini project, but the satellite also constitutes Cassini’s “tour engine”. At a diameter of 5150 km, Titan is the only satellite

large enough to impart a significant change on Cassini’s (and Voyager 1’s) trajectory during a flyby, making it the only means – other than the expenditure of propellant – to allow a spacecraft orbiting Saturn to maneuver quickly around the planetary system. Without Titan, spacecraft would be stuck in a nearly static orbit, slowly precessing due to Saturn’s oblateness and its motion around the Sun. Rhea, Saturn’s second largest satellite, has only 2% of the mass of Titan, offering only a tiny fraction of Titan’s capability to redirect spacecraft orbits. Each Titan flyby can impart a total change in velocity (Δv) of up to 800 meters per second, which is more than adequate for rapid transit around the Saturnian system in the span of months. The total amount of Δv that Cassini had at its disposal after arrival at Saturn was only 560 m/s – less than that provided by a single Titan flyby. (By the end of the Cassini mission, it will have made 127 close flybys.) Therefore, Titan was a prime target not only for scientific reasons, but for engineering and programmatic reasons as well. It is not an exaggeration to say that Titan enabled nearly all of Cassini’s scientific study by providing the bulk of the propulsion required to visit its targets of interest within a reasonable mission lifespan.

2. THE CASSINI ENVIRONMENT PRE-ARRIVAL

The Cassini project carried a number of Saturn trajectory (or “tour”) designs since the start of Phase A in 1990, but Titan flyby planning came into sharp focus halfway through its cruise period in 2001, when the project was conducting the final tour design pre-arrival. A key question that faced the project was evaluating the lowest altitude above Titan’s surface that Cassini could traverse safely. This was of great interest to nearly everyone on the project, given the scientific, engineering, and programmatic value of the flybys. Among Cassini’s scientific instruments was the Ion and Neutral Mass Spectrometer, an instrument specifically designed to study Titan’s atmosphere and its constituents (as well as Saturn’s atmosphere and the ring environment).

A summary of some of the stakeholders and their interests in the selection of these altitudes is given in Table 1.

Table 1. Stakeholders in the Titan flyby altitude selection

Stakeholder	Interests
Project Manager	Project success; risk; cost
Project Scientist	Overall science
Spacecraft Manager	Safety; resources
Trajectory Designers	Trajectory design space
Navigators	Delivery performance
Atmospheric Scientists	Science
Non-atmospheric Scientists	Science (different)
Sequencing Team	Stability; replanning

The Cassini/Huygens hardware consists of the orbiter with a dozen instruments, as shown in Figure 2, and the Huygens Titan Probe with six instruments. The orbiter's science instruments are body fixed and the entire spacecraft must be turned to point them. Consequently, most science observations are made without a real-time communications link to Earth and stored on Cassini's solid state recorder for later playback. Cassini is three-axis stabilized and controls its attitude using either its reaction control system thrusters or reaction wheels. The spacecraft is frequently reoriented as operations requires a variety of attitudes each day, including uplink and downlink that requires the high gain antenna be pointed at Earth. Spacecraft articulation is typically performed with the reaction wheels. Thrusters can impart a significantly higher torque than the reaction wheels, but have lower stability, and so are used rarely: when fast turn rates are required; when orbital maneuvers are performed using either the thrusters or the main engine; or when there is external torque on the spacecraft which exceeds the capabilities of the reaction wheels – such as drag from Titan's atmosphere.




Figure 2 – Cassini Spacecraft Configuration

As of 2001, the Cassini spacecraft was operating nearly flawlessly. It had successfully completed its last flyby at Jupiter in December of 2000 and was on the long leg to Saturn.

Based on the interests listed in Table 1 combined with the environment and operational constraints, one might expect the stakeholders to have the following inherent biases listed in Table 2. Cruise science had been approved and initiated (including at Jupiter) and the science and engineering teams were learning how best to operate the spacecraft at Saturn. The project was working on the Huygens probe anomaly, developing strategies to correct a design flaw which would

prevent the Huygens probe support equipment from receiving probe data with the expected Doppler shift from a high-speed flyby of Titan. The redesign of the Huygens probe delivery trajectory would solve the problem, but this issue was of concern to the project, NASA, and ESA, and constituted a significant risk at the time to mission return.

Table 2. Inherent biases by stakeholders in the Titan flyby altitude selection

Stakeholder	Inherent Bias
Project Manager	Go high; minimize risk
Project Scientist	Go low; maximize science
Spacecraft Manager	Go high; minimize risk and resource usage
Trajectory Designers	Go low; maximize Δv per flyby
Navigators	Go high; maximize acceptable errors to trajectory targets and minimize uncertainties
Atmospheric Scientists	Go low; sample densest atmosphere possible
Non-atmospheric Scientists	Go high; prioritize Titan close approaches for own science
Sequencing Team	Go high; minimize odds of replanning or safe mode "on their watch"

Also in the forefront of the minds of project and NASA personnel was the recent double loss of both Mars Polar Lander and Mars Climate Orbiter in late 1999. Post-mortem analyses had shown the most likely cause of the losses was premature termination of the engine burn for the lander, and errors in navigation of the orbiter traced to a mismatch of units in estimating small forces on the spacecraft on approach to Mars which led to it being directed to a trajectory much lower than expected, leading to disintegration in the atmosphere. The Mars Polar Lander failure mode was not particularly relevant to Cassini Titan flybys, but Mars Climate Orbiter's was directly applicable, and the loss of two missions so recently in JPL's history weighed heavily on all project managers – including Cassini's when faced with similar circumstances in navigating the Saturnian system, especially Titan.

Naturally, many of these inherent biases are at odds with each other, in nearly equal strength. In solving this problem, as with many others the Cassini project confronted, mission planners came to coin the adage: the natural world conspires to make decisions as difficult as possible.

Of course, Tables 1 and 2 are over-simplifications. The project teams were well informed of the project's objectives and NASA's recent history and appreciated the larger goals and concerns of science and mission success, and battle lines were not so bluntly drawn. Nevertheless, discussions of this topic were frequent, often energetic, and echoed all of the positions listed above.

The job of leading the resolution of these biases, and building consensus leading to a decision on Titan minimum flyby altitudes, fell to the Cassini mission planners (the first author being the lead at the time) – the planning-related systems engineers of the project. And what is to be the inherent bias of the systems engineer? What is the natural position a systems engineer in this position should take, at the focus of all of this project attention from so many parties, with the fate of the mission, or at least measurable science return and mission risk, on the line? Which side should they take?

There is only one answer to this question, and of this there can be no debate: it is the job of the systems engineer to have no inherent bias. It is their job not to guess, but rather to rely on data and hard-nosed technical analysis.

3. TITAN FLYBY PLANNING PRE-ARRIVAL

The first step in selecting the minimum altitude for Titan is to determine the driving limitation. The interaction of the atmosphere with the spacecraft has a variety of possible effects, including:

- Alterations to the trajectory from drag force during the flyby
- Mechanical stresses from drag force
- Loss of attitude control from torque imparted due to an offset between the center of mass and center of pressure
- Heating from high speed impact of gas molecules
- Ionization of spacecraft materials from chemical interactions
- Expected spacecraft responses to any of the above that require excessive ground planning to manage
- Unexpected spacecraft response to any of the above that cannot be anticipated (unknown unknowns)

Thorough analyses of all of the above effects had been conducted in early development and operations, and it became clear that the first effect that would cause problems for Cassini (i.e., become a limiting factor at the highest altitude) was loss of attitude control due to torque on the spacecraft. Based on the early atmospheric models, this would be expected to occur somewhere in the 900 – 1000 km range above the surface (see Figure 3). Trajectory effects, mechanical stresses, heating, and ionization are not a problem at these altitudes, and would only be factors lower down. As to the latter two effects above, spacecraft interactions with tenuous atmospheres were well understood at the time, and Titan’s atmosphere of nitrogen, methane, hydrogen, hydrocarbons, and noble gases did not constitute any unusual hazards to spacecraft materials. However,

vigilance to unanticipated spacecraft responses could not be forgotten.




Figure 3 – Titan’s upper atmosphere areas of concern

It is important to note that loss of attitude control is not equivalent to loss of mission. Instead, a safing response is triggered on board the spacecraft which would stop the background sequence in execution – leading to complete loss of science, temporarily – and attempt to recover the spacecraft to a safe altitude as soon as could be managed, as well as initiate direct communication with the Earth to transmit anomaly telemetry and wait for instructions. Typical flyby speeds for Cassini at Titan were expected to be about 6 km/s, and at that speed, the orbiter’s interaction with the atmosphere would last only for a minute or two, with peak atmospheric torque only for a few tens of seconds. Any safing event is a bad day for Cassini, but not a loss of mission – unless modeling or navigation delivery should prove to be extremely wrong (as was the case for Mars Climate Orbiter). The negative impacts, therefore, that are at stake if the minimum flyby altitudes are wrong are loss of science, embarrassment of the project and potential added oversight, and significant replanning and associated labor costs. The loss of science impacts not only Titan atmospheric scientists, but all Titan scientists and other teams as well, since any safing event would likely cause a loss of several days’ worth of science before the background sequence could be restarted. With a flagship mission and 12 scientific instruments at such a target-rich environment as Saturn, several days of lost science is significant and likely to be felt by multiple teams.

To determine the “tumble altitude” or minimum acceptable safe flyby altitude, the problem remained to determine the mass density of Titan’s atmosphere as a function of altitude, as well as the spacecraft characteristics (center of mass and pressure, projected area, thruster performance, etc.) to determine where the atmospheric drag would exceed the orbiter’s capability to counteract it. This is a relatively straightforward problem frequently posed in university level aerospace classrooms, with a solution determined by solution of the following equations:

$$F_{drag} = \frac{1}{2} \rho V^2 C_d A \quad (1)$$

$$t_{drag} = F_{drag}(c.p. - c.m.) \quad (2)$$

$$t_{thrusters} = \sum F_{thrust}(t.l. - c.m.)\cos\theta \quad (3)$$

In Equation 1, the force of drag is a function of the atmospheric density (ρ), the velocity of the spacecraft (V) through the atmosphere, the coefficient of drag (C_d), and the spacecraft's projected area as seen by the atmosphere (A). Equation 2 turns this drag force into a torque by multiplication by the moment arm, i.e., the offset between the spacecraft's center of pressure (c.p.) and center of mass (c.m.). The counteracting torque that can be provided by the thrusters is determined by the sum of the thrusting force from each thruster firing multiplied by its moment arm determined by the distance between the thruster's physical location (t.l.) and the spacecraft center of mass, multiplied also by the cosine of the angle between the direction the thruster torque is imparted and the axis about which the atmospheric torque is being applied (θ). Generally, Cassini was oriented with a principal axis "into the wind", so only a subset of thrusters were used and all of their thrusting was used to counteract drag with angles (θ) equal to zero.

With all of the above parameters except atmospheric density known, all that has to be done is to require that the torque from the thrusters must be higher (by an amount that seems prudent from a risk standpoint) than the drag torque, solve for density, back out the altitude from an atmospheric model, and the problem is solved. Or at least, this is commonly how it is solved at the university level.

Enter the "real world". The reality was that none of the spacecraft characteristics listed in Equations 1-3 were known perfectly. It is not common practice to place fully assembled spacecraft in high-speed atmospheric test chambers to accurately measure the projected area, coefficient of drag, and center of pressure at all possible attitudes under the effects of tenuous gas traveling at 6 km/s, nor did chambers with such capabilities that can accommodate a spacecraft the size of a small school bus exist in the 1990s (and may not to this day). Furthermore, both the center of mass and the thruster performance were expected to vary with time, as fuel was expended and sloshed within the propellant tanks and management devices. The hydrazine system that supplied propellant to the thrusters was a blowdown system, with one pyro recharge of helium pressurant available, but not scheduled, that would make a radical change in the thrust available – nearly doubling it. So the spacecraft characteristics were highly dependent on interpretation and mission progress that was difficult to predict. Problems such as this can be solved by statistical and/or Monte Carlo analysis where each parameter is assigned a statistical range or distribution and results where the spacecraft loses attitude control, or "tumbles" is set to an acceptably low probability.

Regarding an atmospheric model for Titan, knowledge of Titan was limited in 2001, and contributed even more uncertainty to the problem. As previously stated, the

primary source of data for Titan's atmosphere came from the Voyager 1 flyby, and in particular, the occultation of the Sun by Titan's atmosphere observed by Voyager's ultraviolet spectrometer. The first engineering model to be developed from this data came to Cassini in 1987 from Lellouch and Hunten, but this was followed by no fewer than three later models, all of which showed different density profiles by different scientists, most of which were affiliated with the Cassini project. Among the other models, illustrated in Figure 4, was one from Roger Yelle. Roger had been named as the chair of Cassini's Titan Atmospheric Modeling Working Group (TAMWG), whose primary purpose was to deliver (and maintain) just such an engineering model to the project to use for Huygens probe planning and other studies such as orbiter minimum altitude selection.




Figure 4 – Early Titan atmospheric models (1994) – Yelle (Ye) and Lellouch and Hunten (LH)

As sophisticated as the scientific community was in its expertise and analysis of Voyager 1 data, there was still a very large inherent limitation in these atmospheric models: they were all derived from one source, taken at one location above Titan, at one point in time and space. Not only that, but the minimum altitude to which the Voyager UVS data

extended was 1150 km – well above the altitudes that were likely to be the minima for the Cassini orbiter. Roger Yelle himself cautioned the project in a memo in 2001:

“Far too much faith is being placed in these model atmospheres. The models are based on Voyager data and several assumptions that can only be tested by Cassini measurements. These include assumptions about the temperature profile below 1150 km (the lowest altitude measured by the Voyager UVS), the assumption that the atmosphere is as it was during the time of the Voyager encounter, [and] the assumption that the latitude and longitude measured by Voyager are representative of the latitude and longitude through which the Cassini orbiter will pass. We do not understand Titan’s atmosphere well enough to account for changes in time, latitude, and longitude in a quantitative way. Thus, it must be recognized that Titan’s atmosphere will differ from the model predictions and that Cassini will have to fly at different altitudes than are currently planned. We will only know the correct altitude for the low-altitude passes after Cassini instruments have measured the atmospheric density.... All such studies should be based on a range of models, not the assumption that a single model is correct.”

Captured in the second to last sentence of Roger’s quote lies the rub, the classic catch-22 facing mission planners on this issue, and indeed in similar situations in most NASA missions. The Cassini project had to know Titan’s atmosphere to design their plans for exploration. But they didn’t yet know Titan’s atmosphere. That’s why Cassini was going there in the first place. The project had to do its best to make crucial mission planning decisions with limited information. Sometime, that’s spaceflight systems engineering. Those calls just have to be made, and made with equal parts prudence and courage.

Given these concerns, unknowns, and variations in the parameters relevant to the study, the ideal process that should have been adopted by the project for the pre-arrival Titan minimum altitude selection is as follows:

- Gather the best scientific minds on Titan atmospheric modeling (in other words, the TAMWG) and build consensus on a best estimate of the atmospheric density with altitude, as well as the extremes of the reasonable range of variation (i.e., the “range of models”)
- Gather the best engineering minds with Cassini spacecraft expertise, and build consensus on the best estimate of drag characteristics and thruster capabilities of the spacecraft with time, as well as the extremes of the reasonable range of variation
- Study the partials of every parameter that affects minimum altitude selection to determine which ones the altitude is most sensitive to, and therefore deserve the most scrutiny

- Decide upon the risk posture of the project, in other words, how much margin to leave between the expected atmospheric density and the maximum density the spacecraft could handle; this was expressed as a “maximum duty cycle”, in other words, the percentage of the thruster capability that was expected to be used at the lowest flybys (less than 100%; 50% would indicate the thrusters were firing pulses half of the time at closest approach)
- Combine all of the above in a statistical analysis, perhaps employing Bayesian statistics, which provide results on the likelihood of loss of attitude control as a function of Titan flyby altitude
- Expose the results to members of the project and invite review to ensure that the analyses have been assembled correctly
- Select a minimum Titan flyby altitude or a flyby altitude profile with time, as well as a plan to learn from early flybys and adopt updated models into later flybys; study the operational implications of changing flyby altitudes during the tour

Indeed, the project did follow most of the above process. There was much socialization of the analyses in a project-wide forum where all interested parties were welcome. After many discussions and review of models and process described above, the Cassini project selected a maximum duty cycle of 60% for Titan flybys, which elicited the minimum flyby altitude of 950 km above the surface. However, given the high uncertainties in the models, and perhaps influenced by risk aversion and recent events, the project felt that it was wisest to step down to this altitude over several flybys. Therefore, a “toe-dipping” method was put in place to start higher up, assess the navigation capabilities and pre-arrival models, and gradually go deeper into Titan’s atmosphere. The first low Titan flyby would not be immediately at the minimum altitude of 950 km; the plan was to achieve this altitude on a later flyby. The first flyby was planned to be at 1200 km and, leaving adequate time for analysis and avoiding any risk to the Huygens mission on the third flyby, the project would not go to 950 km until the fourth flyby, after the Huygens mission was over. This flyby would take place four months after the 1200 km flyby to give adequate time to analyze and react to flight results.

Note that not all Titan flybys were free to have altitudes selected at the will of the project, e.g., at 950 km. Often there were reasons related to celestial mechanics to fly higher than at the minimum altitude; some times the trajectory designers simply did not need, and could not use, the maximum Δv possible from each flyby. Only about half of the Titan flybys would eventually be set at the minimum altitude.

4. EARLY MISSION RESULTS

The Huygens probe trajectory redesign conducted in the years before arrival replaced the first two flybys – initially called T1 and T2 – with three flybys named Ta, Tb, and Tc, with the Huygens mission itself on Tc. The orbiter flew by Titan at an altitude of 60,000 km for this flyby, but Ta and Tb were less constrained. Ta was selected to be at 1200 km, but Tb could not also be at 1200 km due to geometric constraints for the early tour – at first. The Huygens redesign was to rejoin the original planned trajectory starting with the T3 flyby (which would now be the fourth encounter).

Table 3 shows a summary of Cassini’s early mission flybys. Encounters with altitudes significantly above 950 km are shown in *italics*.

Table 3 – Cassini’s Early Flybys

Titan Flyby	Date	Planned altitude target pre-arrival (km)	Updated altitude target at execution (km)	Altitude as flown (km)
Ta	10/26/2004	1200	1200	1174
Tb	12/13/2004	2197	1200	1192
<i>Tc*</i>	<i>01/14/2005</i>	<i>60000</i>	<i>60000</i>	<i>60003</i>
<i>T3</i>	<i>02/15/2005</i>	<i>1000</i>	<i>1577</i>	<i>1580</i>
<i>T4</i>	<i>03/31/2005</i>	<i>2509</i>	<i>2402</i>	<i>2404</i>
T5	04/16/2005	950	1025	1026
<i>T6</i>	<i>08/22/2005</i>	<i>4007</i>	<i>3660</i>	<i>3660</i>
T7	09/07/2005	950	1075	1075
<i>(8 high altitude flybys)</i>				
T16	07/22/2006	950	950	950
T17	09/07/2006	950	1000	1000

*Huygens mission

The first flyby – Ta – took place in October 2004 at an altitude of 1174 km. This is measurably lower than the target of 1200 km, but given the fact that this was Cassini’s first targeted flyby of Titan and took place 23 years after the previous flyby, this error was well within expected navigation performance and was a clear success. The atmospheric density as reconstructed by the spacecraft team, however, was disturbingly higher than expected. The duty cycle for this flyby was 6%, and based on our models, the T5 and T7 flybys at 950 km would see a drag force 16x greater – in other words, very near the maximum thruster capability at 100% duty cycle. Due to this surprise, all of Cassini’s flybys targeted at 950 km were considered to be at risk of tumbling. This was not a good result.

In parallel with this unexpected atmospheric event, the project identified a potential concern with the orbiter and probe’s passage by Saturn’s satellite Iapetus on the third orbit, after probe separation. This was a non-targeted, somewhat distant flyby, but it was close enough that Iapetus’ gravity could negatively affect the probe mission,

as Iapetus’ mass was not well known at the time. To increase the distance of this non-targeted encounter, and reduce the gravitational uncertainties to a more comfortable level, the project implemented a trajectory update which increased the Iapetus flyby distance. To accomplish this, the Tb flyby design was altered and its altitude was lowered from 2197 km to 1200 km (offering additional atmospheric data). T3 was also changed, and raised to 1577 km, pulling it well above the minimum altitude thought to be risky. Downstream impacts past T3 were slight, however, so the next low flybys – T5 and T7 in the following months – were still of concern.

The Tb navigation delivery was much better than at Ta, drawing from a much increased accuracy in the Titan ephemeris provided by the Ta flyby. Minimal atmospheric data was collected at Tb due to conflicting science requirements, so the project relied solely on Ta data to determine what to do with T5 and T7 and raised them from 950 km to 1025 and 1075 km respectively.

The expected thruster duty cycles at the T5 and T7 flybys, after their altitudes were raised, were 60% and 40% respectively, but the actual result from T5 was 20%, and from T7, 14%. Again, the actual densities did not match the expected densities – attributed, in part, to the extrapolation from results at 1174 km to altitudes much lower. Despite T5’s better than expected results, at 20% duty cycle, T7 was not lowered to 950 km, even though the flyby occurred four and a half months later. The project had already adopted the trajectory update with the combined T5 and T7 changes and was already starting to feel the costs of significant replanning from multiple trajectory updates. With all of Cassini’s instruments body-fixed, science teams were required to work together in often contentious meetings to agree on pointing and data collection, and the replanning of flybys and sequences driven by trajectory updates was taking its toll. The planning process was extremely complicated and time-consuming, and it was better to wait until the next opportunity to realize the first flyby at the minimum altitude.

In the midst of these unexpected density results, there was an additional challenge. Density information was coming primarily from two sources: direct atmospheric measurements from the INMS instrument, and derived measurements from thruster firings as analyzed by the attitude control (AACS) team. The INMS data and AACS data differed consistently by a factor of 3, a mystifying result. While AACS data showed the direct impact of drag on the spacecraft, which was the effect of most concern, INMS directly measured the atmospheric density with altitude, the principal Titan variable of interest in the drag calculations. After significant study by both the AACS and INMS teams early in the tour, no solutions presented themselves, and this discrepancy remained a mystery for years. Since the AACS data directly reflected the actual impact of atmosphere on the spacecraft, which is ultimately the effect that is being managed, those results were left

alone and the INMS data were scaled (up) by a factor of three to align with AACS data. The project has since confirmed, after the end of the prime mission, that this was the right strategy, as updates to the modeling of gas flow into the INMS instrument have resolved the discrepancy; INMS has simply ingested significantly less gas than originally expected (by a factor of ~ 3) based on the physical layout of their instrument.

5. THE T16 ANALYSIS

Following the events of the early tour, the first minimum flyby altitude at 950 km was finally to be at T16, in July of 2006 – nearly two years after Cassini’s arrival at Saturn. Preceding T16 were an array of higher altitude flybys while the orbiter remained in Saturn’s equatorial plane performing, among other things, targeted flybys of Saturn’s other satellites including Hyperion, Dione, and Rhea, with little need for maximum Δv from Titan and therefore, minimum altitude flybys. With meticulous analysis possible over this extended time period, the project and TAMWG began a more systematic study of the few new measurements of atmospheric density available with respect to a variety of parameters to determine if trends could be observed that should be incorporated to predict future performance. One such plot caught the attention of the project, and is shown in Figure 5.




Figure 5 – Density variation with latitude, from INMS, AACS, and HASI (Huygens) measurements. All densities are normalized to 1450 km altitude for comparison.

This plot illustrates the density measurements from Ta, T5, and T7 as a function of the Titan latitude at which they were made. Note that this data contains measurements made at a variety of altitudes, and each measurement is scaled to one single altitude (1450 km, arbitrarily) using a relative density-with-altitude model (which was admittedly imperfect) so that direct comparisons between measurements could be made. This plot also scaled up INMS data by a factor of three, as previously discussed. A Huygens data point is also shown from HASI (the probe’s Huygens Atmospheric Structure Instrument); however,

since the discrepancy between AACS and INMS measurements had not yet been resolved, the project was uncertain as to whether the factor of 3 should be applied to that point as well – therefore two points are shown with a line joining them.

Based on Figure 5, it is plausible to conclude that there is a trend of higher densities found near the equator and lower ones near the poles. The grouping of density points in the northern hemisphere was interpreted as particularly convincing. The TAMWG generated a variety of curve fits that would emulate the behavior observed in Figure 5, and these fits are illustrated in Figure 6. After scrutiny and discussion, the solid line in this Figure was selected for subsequent Titan atmospheric modeling with respect to latitude. This line is symmetric about the equator.

It was at this point that the project departed from its earlier strategies and committed a handful of minor mistakes that would prove, fortunately, to be of little consequence to the project. As an exercise before continuing, the reader could try and identify four distinct problems with this most recent approach leading to the solid-line latitudinal model in Figure 6.




Figure 6 – Density variation with latitude, as in Figure 4, with fitted models

First, it should be recognized that it is human nature to fit smooth lines through data points, and the data collected to date is really quite insufficient to support the conclusion that one simple model is likely to be accurate here. In Roger Yelle’s own words from earlier in this paper, “studies should be based on a range of models, not the assumption that one single model is correct”. So the selection of one latitudinal model, rather than a range of possible models with piecewise statistical likelihoods, is ill advised.

Second, it is also human nature to look for symmetry in nature where none is guaranteed to exist. The fact that the trend is symmetrical about the equator flies directly in the face of the very earliest close observations of Titan. Figure 1 of this paper – one of the first close-up pictures of Titan –

shows a hemispherical difference, which HST later confirmed, though in the opposite direction. There is no reason whatsoever to assume that Titan's upper atmosphere is symmetric about the equator.

Third, the tight clustering of the data points showing a ramping trend in the northern hemisphere which support the solid-line model is arguably artificial, since it is composed of both AACS and INMS data, the latter of which were artificially adjusted (at the time) to match the former. That should be a factor which erodes confidence that the latitudinal trend is real.

Fourth, the solid-line trend has a nonzero slope at the poles. This seems reasonable graphically based on the way the data was plotted, with the poles at the edges of the plot. However, the poles are not a boundary (in the manner of a flat Earth, or Titan); the atmosphere "keeps going" over each pole and down the other side, and this slope would create an awkward inflection point that makes very little physical sense save for a dramatic, permanent low pressure cyclone or similar disturbance at both poles. No evidence of such features had been observed at Titan.

This last problem was especially egregious for the T16 flyby, the closest approach of which was to take place very close to the north pole at 85 degrees north latitude. Based on the solid-line curve's density variation with latitude profile and adjustments made to the pre-arrival density variation with altitude model and uncertainties, the expected thruster cycle for T16 was believed to be only 22%. Therefore, the project elected to proceed with T16 at 950 km as planned with little hesitation. In fact, the project used the T5 and T7 data to redesign the complete suite of remaining Titan flybys in March of 2006, incorporating the individual spacecraft attitudes, thruster usage, and science plans at closest approach for the first time, to fine-tune each flyby to maintain a fixed level of thruster capability margin. This re-tuning resulted in flyby altitudes for the remainder of the tour between 950 km (including T16) and 1030 km.

When T16 occurred several months later, once again the results were not as expected. The solid-line curve from Figure 5 plummeted so severely at the poles that this profile under-predicted the duty cycle significantly, and the actual thruster duty cycle experienced was 62% - nearly three times higher, though not so high as to be a big concern for the project in and of itself. The data and models were clearly not yet comprehensive enough to cover all of the possible variations, and needed further updating to make accurate predictions in the future.

The next low flyby, T17, was to be only 47 days later at an altitude of 1000 km (thankfully not as low as T16, due to the replan earlier that year). The models were updated with the latest data post-T16, and a new latitudinal trend was developed which flattened at the poles, making more intuitive physical sense. Figure 7 shows the updated model as the blue line. This blue curve corrected the counter-intuitive inflection at the poles, but admittedly was still

symmetric about the equator. For T17 assessment, a line was drawn through the +23 degree latitude position, the closest approach latitude for T17. The maximum acceptable density at 1000 km equated to a point on this plot at 8.5×10^{-12} kg/m³, shown as the high blue point on the line. It seemed clear that regardless of what latitudinal variation is assumed, the tumble density would not likely be reached at T17 as no curves or data points (save for the inflated Huygens measurement) ventured anywhere near this point. Therefore, the project elected to continue with T17 as planned. This trend continued throughout 2006 and beyond as data was added and analyzed by TAMWG in support of Titan atmospheric modeling.




Figure 7 – Density variation with latitude with updated T16 and other data and T17 predictions

6. THE REST OF THE MISSION

T17 and subsequent flybys were successful, with no flyby coming much closer than T16 to "maxing out" the thruster performance capabilities – i.e., requiring a 100% (or greater) duty cycle. Only one other flyby – T57, at 955 km and 42 degrees south latitude – would break the 60% duty cycle barrier, at 69% duty cycle in June of 2009. Nearly all of the flybys remained within the range of 25-55% duty cycle. Over time, the project gained confidence that the models in use plus the margin between the thrusters' capabilities and the atmospheric torque were adequate to keep the spacecraft from a safing response. The science community was satisfied with the Titan atmospheric science. Though further tour trajectory adjustments were made, they were minor and the sequencing teams and project managers were satisfied that none were driven by altitude adjustments due to Titan atmospheric variations or updates in modeling. The reaction control system was recharged per plan, giving the hydrazine system new life on the predicted schedule.

The Cassini project has continued to analyze the Titan atmospheric data, searching for trends with respect to a variety of environmental conditions. For a time, densities seemed to decrease over time, as shown in Figure 8.




Figure 8 – Density variation over the years, normalized to 1000 km, with solar activity

It was speculated that this trend could be attributed to solar activity, also shown in Figure 8 in the form of the 10.7 cm solar radio flux (green line). However, this trend was at least obfuscated by the fact that early mission flybys had closest approaches which were predominantly in Titan's northern hemisphere, with later flybys in 2008-2010 in the southern hemisphere. Furthermore, Saturn's increasing distance from the Sun was shown to have more influence on solar heating of Titan's upper atmosphere than the solar cycle.

It was later speculated that mass densities might be connected with longitude, since densities measured in the southern hemisphere were on average 30% lower, for a time. These lower densities fell mostly in to the 60° - 180° longitude sector implying possible magnetospheric forcing.

Years continued to pass, and geometric-related trends remained a possibility, but correlation has not proven causation, and there are still not enough measurements to prove one way or the other whether Titan's upper atmosphere varies randomly, or is tied to one or a set of environmental conditions. The plot of density versus latitude is now more reminiscent of a random scatter plot, possibly with hemispherical asymmetry with higher density in the northern hemisphere, as shown in Figure 9. Fortunately, the models implemented in Titan flyby design in 2006, and in subsequent years for the first and second mission extensions, have continued to serve the project well with no need for redesign.




Figure 9 – Titan atmospheric density normalized to 1000 km; AACS data in triangles and INMS data in circles

7. CONCLUSIONS

As of 2017, the Cassini team has a much better understanding of both Titan and the spacecraft. They now understand the discrepancies between INMS and AACS; the INMS instrument was under-reporting density based on mis-modeling of the flow of relatively high pressure gas within the instrument (a factor which did not affect the instrument's prime science goals or performance). The spacecraft has also been in operation around Saturn, studying Titan, for more than a decade. Its atmospheric fluctuations have not resulted in drag large enough to risk loss of attitude, nor small enough to prevent groundbreaking science from being collected. Too many of the parameters are intermingled, and not sampled with great enough spatial and temporal frequency to fully resolve this question. It will likely take decades, or measurements spaced across decades, and therefore future missions, to resolve the question of Titan's atmospheric variation.

Now that the team is nearing the end of the mission they look towards the last, and very important, low flyby – T126, Cassini's 127th flyby of Titan, on April 22, 2017. This final flyby will single-handedly set up the ballistic trajectory for Cassini's Grand Finale, when the spacecraft will pass 22 times between the rings and the planet, and thereafter into Saturn on September 15th 2017. These orbits are equivalent to a whole new mission for Cassini, one never envisioned before 2008, and offers a suite of new and unique science enabled by its Juno-like close passages by Saturn and the rings. It is crucial that this flyby go as planned. The targeted altitude for the T126 flyby is 979 km, with a latitude of +66 degrees at closest approach. It is in family with other flybys and in an area of the atmosphere visited many times. However, Titan has surprised us before, and five months will have passed since the last encounter. Titan may yet have some tricks remaining.

The Cassini mission planners and the rest of the project team had a challenging task before them, starting in 2001.

Some key takeaways from this chronicle of systems engineering are listed below.

- Statistical analysis must be done carefully, especially with limited data sets. The urge to apply human intuition to derive trends from limited statistics is often strong. Don't guess – rely on data. Knowledge of statistical techniques and distributions is crucial, especially Bayesian analysis. Not all distributions are symmetric; nature is not always Gaussian.
- Key project trades are rarely ever finished. They must be continually reassessed with new information and new perspectives. Room must be left for new data to inform and update trades over time, and a graduated risk posture, reducing margin with time as knowledge improves, is worth consideration.
- The interests and inherent biases of all stakeholders must be understood. Biases can be countered with technical analysis to build consensus.
- With high uncertainties comes the need to build a range of models that span the extremes of the reasonable range of interpretation. All plausible models should be treated equally.
- Care must be taken when plotting data to illuminate the limitations of the plots.
- Systems engineers should embrace uncertainty and be comfortable with chaos and change.

ACKNOWLEDGEMENTS

This work was performed at the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration.

The authors wish to acknowledge the following personnel who have contributed so much to this field over the course of the Cassini project lifespan.

- Cassini mission planners (not already authors): Emily Manor-Chapman, Erick Sturm, James Frautnick, Gene Y. Lee
- Cassini spacecraft managers: Julie L. Webster, Earl Maize
- Cassini project managers: Robert T. Mitchell, Earl Maize
- Roger Yelle, Titan Atmospheric Modeling Working Group lead
- Contributing scientists to Titan atmospheric modeling and interpretation: J. Hunter Waite (INMS PI), Darrell Strobel, Ingo Mueller-Wodarg.

- Cassini attitude control team leads and members: Allan Lee, Thomas Burk, Luis Andrade, Sam Sarani
- Cassini trajectory designers: Brent Buffington, John Smith, Nathan Strange
- Cassini navigation team leads and members: Jeremy Jones, Duane Roth, Fred Pelletier, Yungsun Hahn

REFERENCES

- [1] *Engineering Models for Titan's Atmosphere*, R. V. Yelle, D. F. Strobel, E. Lellouch, and D. Gautier, In: HUYGENS Science, Payload and Mission, ESA SP-177, 1997.
- [2] *Titan atmosphere engineering model*, Lellouch, E. & Hunten, D. M. ESA Space Science Department Internal Publication, ESLAB 87-199, 1987.
- [3] *Titan's upper atmosphere: Structure and ultraviolet emissions*. Strobel, D. F., Summers, M. E. & Zhu, X. Icarus 100, 512-526, 1992.
- [4] *The Vertical Structure of Titan's Upper Atmosphere from Cassini Ion Neutral Mass Spectrometer Measurements*, R. V. Yelle, N. Borggren, V. de la Haye, W. T. Kasprzak, H. B. Niemann, I. Mueller-Wodarg and J. H. Waite, Jr. Icarus, 182:567-576, 2006.
- [5] *Non-LTE models of Titan's upper atmosphere*, Yelle, R. V., Astrophys. J. 383, 380–400, 1991.
- [6] *NASA Systems Engineering Handbook*, SP-2007-6105 Rev. 1, 12/2007
- [7] *Twenty Years of Systems Engineering on the Cassini-Huygens Mission*, Emily Manor-Chapman, 2013 IEEE Aerospace Conference, 2013.
- [8] *The Cassini-Huygens Mission: Volume 1: Overview, Objectives, and Huygens*, ed. C. Russell.
- [9] "Lessons Learned" excerpt from *Loki Lego Launcher 1 project plan*, Rebecca and Kimberly Yeung, September 2015. <http://www.geekwire.com/2015/interview-seattle-girls-launch-balloon-to-the-edge-of-space-and-nasa-takes-note/>
- [10] *Titan Model Atmospheres and Tour Planning*, R. Yelle, Cassini internal memo, 2000.
- [11] *The Dragons of Saturn: An Insider's Chronicle of the Engineering of NASA's most successful planetary mission*, D. Seal, 2017.
- [12] *Titan*, in *Saturn* (eds. T. Gehrels, M. S. Matthews), Hunten, D. M., M. G. Tomasko, F. M. Flasar, R. E. Samuelson, D. F. Strobel, and D. J. Stevenson, Tucson: University of Arizona Press, pp. 671-759, 1984.

BIOGRAPHIES



David Seal is JPL's principal engineer for mission planning, and has been working in the field since departing MIT where he received his B.S. / M.S. in Aerospace Engineering in 1990 / 1991. He has been the lead planner for the Cassini/Huygens mission, the

Shuttle Radar Topography Mission (STS-99), and the NASA/ISRO SAR mission (NISAR). He currently supports JPL's Mission Systems Engineering section, JPL's planetary mission formulation program office, and NASA's Agency-wide Mission Operations Capability team in the domain of operations planning.



Molly Bittner is a Systems Engineer in the Autonomy and Fault Protection Group of the Flight Systems Engineering, Integration and Test Section (313) at the Jet Propulsion Laboratory (JPL). Currently, she is a Systems Engineer on the Cassini Spacecraft Operations

(SCO) team, leading both engineering sequence development and OTM design and execution. She is also a Mission Planner, and recently published "The Last Orbit: Planning Cassini's Plummet into Saturn", along with the Cassini Mission Planning Lead, Erick Sturm. Bittner has worked on several other projects including: VASE, Rov-E and Mars 2020. Bittner earned a B.S. in Aerospace Engineering from the Georgia Institute of Technology. She is currently pursuing an M.S. in Astronautical Engineering from the University of Southern California with a focus on Systems Engineering.